

UNCLASSIFIED

REPORT DOCUMENTATION PAGE

AD-A212 121

8 1989

1b. RESTRICTIVE MARKINGS

3. DISTRIBUTION/AVAILABILITY OF REPORT

Approved for public release; distribution unlimited

4. PERFORMING ORGANIZATION REPORT NUMBER(S)

5. MONITORING ORGANIZATION REPORT NUMBER(S)

AFOSR-TR-89-1222

6a. NAME OF PERFORMING ORGANIZATION

Department of Psychology

5b. OFFICE SYMBOL
(if applicable)

7a. NAME OF MONITORING ORGANIZATION

Air Force Office of Scientific Research (NL)

6c. ADDRESS (City, State, and ZIP Code)

University of California
Berkeley, CA 94720

7b. ADDRESS (City, State, and ZIP Code)

Bolling Air Force Base, D.C. 20332-6448

8a. NAME OF FUNDING/SPONSORING
ORGANIZATION

AFOSR

8b. OFFICE SYMBOL
(if applicable)

NL

9. PROCUREMENT INSTRUMENT IDENTIFICATION NUMBER

AFOSR-89-0206

8c. ADDRESS (City, State, and ZIP Code)

Building 410

Bolling AFB, DC 20332-6447

10. SOURCE OF FUNDING NUMBERS

PROGRAM
ELEMENT NO.

61102F

PROJECT
NO.

2313

TASK
NO.

A4

WORK UNIT
ACCESSION NO.

11. TITLE (Include Security Classification)

Norms and the perception of events.

12. PERSONAL AUTHOR(S)

Kahneman, Daniel

13a. TYPE OF REPORT

First Annual Technical

13b. TIME COVERED

FROM 6-1-88 TO 7-15-89

14. DATE OF REPORT (Year, Month, Day)

July 28, 1989

15. PAGE COUNT

40

16. SUPPLEMENTARY NOTATION

17. COSATI CODES

FIELD

GROUP

SUB-GROUP

05

09

18. SUBJECT TERMS (Continue on reverse if necessary and identify by block number)

Perception, Categorization, Comparison Processes, Normality
(SDUJ)

19. ABSTRACT (Continue on reverse if necessary and identify by block number)

Much of the first year of the project was spent in developing experimental techniques for the study of the implications of norm theory. The particular focus of the research was the idea of attribute mutability. Several studies were carried out using the dependent variable of self-rated surprise to distinguish surprise governed by absolute novelty and surprise as a reaction to low conditional probability. A major series of experiments examines the roles of attribute mutability and the relations between attributes in the perceived normality of visual objects, agents, and events. Another line of research looks at the role of attention in determining the topic and referent in comparison. New projects examine the felicity conditions for counterfactual language and the effects of extreme category members on category representation. Keywords:

20. DISTRIBUTION/AVAILABILITY OF ABSTRACT

☐ UNCLASSIFIED/UNLIMITED ☒ SAME AS RPT. ☐ DTIC USERS

21. ABSTRACT SECURITY CLASSIFICATION

Unclassified

22a. NAME OF RESPONSIBLE INDIVIDUAL

Alfred R. Fregly, Ph.D.

22b. TELEPHONE (Include Area Code)

(202) 767-5021

22c. OFFICE SYMBOL

NL

NORMS AND THE PERCEPTION OF EVENTS
Report for the period 6-1-88/7-15-89
Daniel Kahneman

Overview

The main objective of the proposed research was to develop and examine the implications of norm theory (Kahneman & Miller, 1986) in several domains. The following theoretical ideas guided the work:

The events and stimuli of experience evoke their own sets of alternatives, to which they are compared. The alternatives can be past experiences retrieved from memory or counterfactual possibilities constructed in imagination. The activation of an evoked comparison set can be viewed as an act of spontaneous categorization. The members of an evoked set tend to share some properties (labeled robust, or immutable) with the evoking event. Other properties (the mutable ones) do not constrain the evoking process, and are therefore allowed to vary freely within the evoked set. The robust properties are used to classify, categorize and assign identities. The mutable properties are used to compare and distinguish events from their natural alternatives. A particularly informative response is surprise, which occurs when an event strongly evokes a very different alternative.

This framework raises a number of questions for research. The following are among the most urgent: what are the rules of the evoking process? which attributes are robust, which are mutable? These questions were the focus of the proposal, and surprise was suggested as a useful dependent variable to study the rules of spontaneous categorization and the relative mutability of different attributes.

Mutability can be characterized both within and between attributes. Some values of an attribute are more mutable than others: for example, we argued that abnormal values tend to be mutable and evoke their normal alternatives, whereas normal values are more robust and do not evoke abnormal alternatives. A more difficult question concerns mutability as a characteristic of the attributes that describe events and objects: are some attributes generally more mutable than others? In the analysis proposed in norm theory the robust attributes control the spontaneous categorization of events, but the mutable attributes control the comparison of a stimulus event to its alternatives and the consequences of such comparisons, notably surprise.

Several hypotheses about the relative robustness or mutability of different attributes were proposed: the initial events in a sequence should be more robust than subsequent ones; shape should be more robust than other attributes of visual objects; characteristics of agents should be more robust than characteristics of actions (Gentner & French, 1988). The question of spontaneous categorization can also

be raised for stimuli that consist of separable objects or parts: we hypothesized that variables such as size or codability would control spontaneous categorization. (As will be shown below, however, this hypothesis did not fare very well).

The questions raised in the proposal did not fit neatly into any established paradigm, and a considerable effort was therefore devoted to the development of different dependent variables. Suzanne O'Curry carried out several experiments using rated surprise as a dependent variable. The studies involving objects in motion generally confirmed our analysis of surprise, and supported the hypotheses that agent identity is a preferred basis of spontaneous categorization of events in which an activity is observed. This work is described in Part 1 of this report.

Most of our time and effort was spent in a large number of experiments using normality judgments as the dependent variable. Suzanne O'Curry and Carol Varey collaborated in the studies of normality judgments, described in Part 2. The studies revealed large and consistent differences in the robustness or mutability of different attributes of objects and events. They also suggested an important distinction between the categorization of an event (a conjunction of features) by its robust attributes and the categorization of a pair of separable objects by one of its constituents. Spontaneous categorization of objects and events readily occurred, but there was little willingness to categorize a compound display by its more salient part, when it was incongruously paired with an object associated with another display.

Some of our work was concerned with the possibility that spontaneous categorization is dominated by the attribute on which alternatives are most easily discriminated. We found several cases in which small differences in shape dominated obvious differences in size or color. In general, the hierarchy of attributes was not very sensitive to manipulations of within-attribute variability. Garner (1974) had made similar observations in his studies of free classification, and we agree with his conclusion that dimensional preferences are a genuine effect. However, some of our results suggest that concurrent variation of several attributes has particularly potent effects. We now hypothesize that spontaneous categorization is controlled preferentially by distinctive conjunctions of properties.

Several other topics related to norm theory were studied. Maria Stone and I have been concerned with the allocation of the roles of subject and referent (or norm) when two simultaneously presented items are to be compared, e.g., to decide which of two lines is longer or whether two visual objects are the same or different. The study exploits asymmetric effects of complexity, predicted by Tversky's (1977) contrast model to identify the stimulus that subjects pick as a topic (see also Agostinelli et al, 1988). The results suggests that the current focus of attention is picked to be the topic, and that the stimulus most recently presented tends to be compared to one presented earlier. The experiments on topic selection are summarized in Part 3.



Availability Codes	
Dist	Avail and/or Special
A-1	

The experiments listed so far were concerned with relatively simple events presented on a screen, but questions about the psychological reality of counterfactual alternatives to reality are perhaps more easily studied by other means. With Carol Varey, I have examined the felicity conditions for the use of expressions denoting salient counterfactual possibilities, such as 'X almost happened' or 'X came close to happening'. The linguistic analysis turned up some interesting results about the psychological reality of (some) counterfactuals, the very precise way in which causal assumptions are indicated by counterfactual statements, the role of emotional involvement, and the asymmetric treatment of normal and abnormal cases in conversations about counterfactuals. I summarize our findings in Part 4.

Two new lines of research are currently under development. Frances Kuo has constructed a set of materials in which unusual members of familiar categories (e.g., shark as a fish, Amish as a religion) are presented, in the context of an imagery task. We hope to trace the effects of a single encounter with an unusual instance on the internal representation of the category. In accord with the currently dominant views that allow hybrid models of category representation -- using exemplars along with abstracted information about the category -- we expect that the unusual instance will selectively affect some measures of the category representation but not others. The guiding hypothesis of this work is that questions that require explicit retrieval of information about the category will be relatively immune to effects of single instances. Questions that evoke less reflective responses (e.g., affectively loaded attitudes, or metaphorical uses) may be more sensitive.

With Maria Stone, I have begun a study of the hypothesis that the similarity of visual objects reflects a mental operation in which one of the objects (the topic) is transformed into the other. Our idea is to expose observers to a large number of occurrences in which a visual object A is seen changing into a different object B, and to test the effect of such perceived transformations on later confusions (and speed in same/different judgments) and on the judged similarity of A to B and of B to A. To test our hypothesis about the special role of perceived transformation, we shall also use similar conditions of successive exposure of the two objects in which there is no perception of apparent motion or transformation.

Part 1. Studies of surprise (Research conducted by Sue O'Curry)

The prevalent view of surprise has been that it is a reaction to novelty or violated expectation (Berlyne, 1960, Patterson, 1930). Norm theory (Kahneman & Miller, 1986) suggests that surprise may be linked to processes of comparison and reminding: an event is surprising if some of its salient properties are abnormal (have much

lower frequency than some alternatives) within the comparison set that the event evokes. The comparison set consists mainly of events that share the robust (immutable) properties of the evoking event. Thus, the robust properties govern spontaneous categorization, and the mutable properties govern surprise, conditional on this categorization.

Three experiments were run with ratings of surprise as the dependent measure. The first study tested the hypothesis that object-action pairs are more likely to be categorized by the agent than by the action. In the second study, the action could be identified before the identity of the agent was revealed. The last experiment explored the relative mutability of some properties of static objects.

Experiment 1

According to norm theory, an event consisting of an object in motion should evoke a norm based on the object, because the event is more easily identified by the agent of action than by the action. Gentner (198?) found that when asked to adjust anomalous noun-verb pairs so that they made sense, subjects consistently changed the verb, suggesting that encoding was based on the noun. In our case, the combination of a novel action property with a familiar object should elicit more surprise than the combination of a familiar action with a novel object or a completely new object/action pair.

Method

Twenty-six Berkeley undergraduates participated in a session lasting 20-30 minutes. Some received course credit, others were paid \$3.00.

The experiment was presented to subjects on a Mitsubishi color monitor controlled by an IBM AT computer. Responses were entered on the numeric keypad and were collected into a data file automatically.

Stimuli were colored letters moving in a characteristic manner to six different locations on the monitor. Each movement lasted 3-4 seconds. By varying speed, pauses, and curvilinearity, fifteen distinctive motion "families" were created. Each motion family contained six members, one for each destination. A different assignment of letters to motions, and of letter-motion pairs to trial types was generated for each subject.

The experiment consisted of three blocks of 48 trials, divided into four subblocks of 12 trials. There were six types of trials, which we labeled by the letters A-F. Each type is defined by a particular pairing of an object with a mode of motion (destination varied randomly). Types A, B, and C were "normal": each appeared four times in the initial 12 trials of a block, and three times in each of the following sub-blocks of twelve trials. Types D-F were novel and remained rare: these events did not occur at all in the first sub-block of 12 trials, and occurred once in each of the next three sub-blocks. In Type D a new letter was seen in the motion usually

associated with type A. In Type E the letter from C trials was paired with a new motion. Finally, Type F consisted of a completely new letter-motion pair. Order within each subblock was randomized with the constraint that no more than two test trials could be shown in succession.

After every trial, a scale appeared on the screen asking subjects to rate how surprising or unsurprising they had found the trial. The scale ranged from 0, "not surprising" to 9, "very surprising."

There were two conditions in the experiment, distinguished by whether subjects saw a sample of the letters or a sample of the motions in the instructions. This was done as a check for any biasing effect of seeing objects or motions first. This variation had no effect.

Subjects were seated in front of the monitor and given brief verbal instructions about use of the numeric keypad and about advancing through the screens of instructions. The instructions included the following sentence: "When we ask you to rate surprise, we want you to rate how surprising what you see is in the context of the experiment, not in the context of what you might ordinarily find surprising in real life."

Results and Discussion

Table 1-1 shows the mean ratings of surprise for the different types of trials on successive sub-blocks. The A, B, and C (old character-old motion) trials were averaged together because there were no significant differences between them.

At the initial presentation, the new-new combination was significantly more surprising than the new-old type [$t(25) = 2.99$], but not significantly different from the old-new type [$t = 0.84$]. By the third presentation, however, the new-new type was significantly less surprising than the old-new combination [$t = 2.92$], and did not differ significantly from the new-old type [$t = 1.35$].

The same trends were confirmed by an analysis of the changes in surprise between the first and the third exposure of the three test types. The rating of surprise did not decline as much for old letter - new motion trials with repeated presentation as for the other test trials. The amount of decline in surprise ratings between the first and third presentations was significantly greater for the new object - new motion type than for either of the other test types [$t(25) = 3.19$, $p < .01$]. The amount of decline was not significantly different between the old object - new motion and new object - old motion trials, [$t(25) = 1.13$, n.s.], but the old object - new motion pairs were rated significantly higher at all presentations.

The pattern of results conformed to theoretical expectations. Surprise differed widely among events of equal experienced frequency. Novelty was important, but not the sole determinant of surprise. A

new behavior from an old character was as surprising as a conjunction of novelty in both aspects, and eventually became more surprising. The predicted difference between the effects of introducing a new agent or a new motion was obtained. In general, then, the results confirm the expected difference in mutability between the properties of actions and of objects (see also Gentner & French, 1988).

Table 1-1
Mean Surprise Ratings by Trial Type and Block
Experiment 1

Trial Type	Block			
	1	2	3	4
Mean of A, B, & C (norm)	3.38	2.34	2.07	1.92
D - New Letter - Old Motion	-	4.08	3.47	2.77
E - Old Letter - New Motion	-	5.24	4.94	4.24
F - New Letter - New Motion	-	5.04	3.86	3.15

Experiment 2

A possible objection to the first experiment is that the asymmetry in the effects of objects and motions arises because subjects first identified the object, then predicted its motion. In this view, surprise arises from specific disconfirmed expectations. To meet this objection, the second experiment presented the same stimuli as the first experiment, but the identity of the letter was masked by a small grey square for the first half of every trial, so that information about the motion became available before information about the moving object. If surprise arises mainly from violated expectations, then the highest surprise should be to a familiar motion paired with a new letter. If reminding is important in surprise, the old letter - new motion conjunction should be the most surprising, because the old letter will evoke a norm that associates a different motion with the letter.

Method

Apparatus, design and procedure were identical to the first experiment, except that subjects had 10 practice trials instead of four. The stimuli were the same as in the first experiment, except that the identity of the letter was masked by a small grey square for the first half of each trial (about 1.5 sec), long enough for the motion to be unambiguously identified.

Results

Table 1-2 shows mean scores for the different types of trials. Again, the three types of norm trials are averaged together. If the expectation-based approach to surprise was correct, the trials that combined a new letter with an old motion should have been the most surprising. This was not the case. On the first presentation,

this type of event was rated significantly less surprising than the other two types [$t(29) = 2.12$, $p < .05$, and $t(29) = 3.92$, $p < .001$, respectively, for the comparisons with the old-new and new-new events]. The latter two events did not differ significantly from each other. There were no significant differences between the conditions on the second and third presentations. As in the first experiment, the trials with the old letter and new motion were rated highest after three appearances, although the difference was not significant in this experiment. Again, the amount of decline of surprisingness was higher for the new letter - new motion trials than for the other types of test trials [$t(29) = 3.10$, $p < .001$].

Table 1-2
Mean Surprise Ratings by Trial Type and Block
Experiment 2, "Occluded" Motion

Trial Type	Block			
	1	2	3	4
Mean of A,B, & C (norm)	3.28	1.90	1.54	1.23
D - New Letter - Old Motion	-	4.26	3.61	3.49
E - Old Letter - New Motion	-	4.96	3.91	3.70
F - New Letter - New Motion	-	5.27	4.14	3.39

General Discussion

Perhaps the most significant result of the first experiment was that the conjunction of a familiar object with a novel action was perceived as more surprising than a completely novel object/action pair. While the role of novelty in surprise is not to be discounted, the results indicated that it is not the only important factor -- nor is rarity. Perception of normality appears to reflect conditional frequency as well as absolute frequency, and the judgments of an event appear more likely to be conditioned on the identity of the agent than on the character of the action. The second experiment addressed the possible role of violated expectations in surprise. Subjects appeared to condition actions on agents even when the action was revealed earlier than the identity of the agent.

We conducted a third experiment, which need not be reported in detail. This experiment employed the same experimental design to assess the relations among four attributes of stationary displays: shape, color, position and number. The results provided clear evidence of the dominance of shape over attributes, except for the comparison of shape and color, which was ambiguous. These findings are in agreement with the conclusions of old work by Heidbreder (Heidbreder, 1948; Heidbreder & Overstreet, 1948) and with more recent studies, e.g., Landau, Smith and Jones (1988). Our results were weak, however, and the subjects found the experiment tedious and artificial. This led us

to explore alternative techniques for the measurement of spontaneous categorization and hierarchical coding.

Part 2

Normality ratings: in search of asymmetries of coding

Much of our experimental work this year has used what we call the normality paradigm, a technique which provides a somewhat unusual measure of the processes of encoding and spontaneous categorization of objects and events. The stimuli in our experiments are static images or short sequences of motion, presented on a CONRAC Model 7250, controlled by an IBM-AT, equipped with an Artist-plus graphics board and run on software developed in our lab.

The set of stimuli used in an experiment is defined by pairs of values on three attributes. Two training patterns (or norms) are defined by distinctive values on these attributes. The remaining six possible combinations of the attribute values define test patterns. For example the two norms may be specified by values of shape (circle/triangle), color (pink/blue) and position (left/right of screen). Another example would be gender (face of boy/girl), expression (frown/smile) and label (nonsense syllables presented on the screen, e.g. UZU/GAR). Subjects are asked to observe the training patterns, which are repeated with high frequency. They answer questions about the normality of an attribute on each test pattern. For each test pattern we ask two questions, on different trials, probing the two features that it shares with its nearest neighbor among the two norms. The dependent variable is the proportion of subjects who rate a particular attribute of a test trial as normal.

An example of a complete experiment is shown in Figure 1, in which the norms are Dark/Circle/Left and Light/Triangle/right. One of the test patterns is a dark triangle on the left of the screen. The nearest norm (reached by changing one feature) is the dark circle on the left. The two questions that will be asked for this test pattern are 'IS COLOR NORMAL?' and 'IS POSITION NORMAL?'.

Note that each of the two attributes that are not queried on a particular trial 'votes' for a yes or no answer to the question. In each case, the vote is split. (In the example above we do not ask 'IS THE SHAPE NORMAL?' because the other two attributes of the test pattern both vote 'No'). The balance of answers provides a measure of the relative weights of the two attributes in determining the response.

The sequence of events in an experiment is as follows. Subjects are first exposed to 12 observation trials in which the two norms occur with equal frequency. They are then shown one of the test patterns and are asked to make a dichotomous judgment of the normality of one of its attributes, by writing Y or N in an answer sheet. This is followed by four additional exposures of the training patterns (two for each), then another test pattern, then four more training trials,

and so on until 12 test trials have been presented. In summary, the experiment consists of 116 trials: each of the two norms is presented 52 times, and each of the 6 test patterns is presented twice, paired on each occasion with a different question. The duration of such an experiment is approximately 4 minutes. We have run about 100 of these experiments so far, with 12-20 subjects in each.

Subjects were run in groups of up to four, and in most of our experiments four such groups were run. A session lasting 45 minutes could include up to ten separate experiments, using unrelated norms. The conjunctions of attributes defining the two norms were different for the different sub-groups participating in an experiment.

What do normality ratings mean?

A natural way to interpret a question about the normality of an attribute is by expanding it. As it stands, the question is ambiguous because it can reasonably be expanded in several ways. Thus, the question 'IS THE COLOR NORMAL' could be understood as an abbreviation of 'IS THE COLOR NORMAL FOR THIS POSITION?' or 'IS THE COLOR NORMAL FOR THIS SHAPE?' or perhaps 'IS THE COLOR NORMAL FOR THIS CONJUNCTION OF SHAPE AND POSITION?'. The correct answers would vary accordingly. In practice, we allow only yes and no as answers to the questions, and hope to infer from the answers how the question was interpreted. The interpretation of the question, in turn, is expected to provide an indication of the role of different attributes in the spontaneous categorization (encoding) of patterns and events: in choosing to evaluate the normality of the position of the stimulus 'for' the shape, rather than for its color, the subject implicitly categorizes the event in one way rather than another.

In most of our research in this general design, position was one of the attributes varied throughout the series of experiments included in a session: in each of the experiments four of the stimuli were presented on the left of the screen, four on the right. In another study one of two nonsense-word labels was presented with each stimulus. We refer to such an attribute (position or label) as a medium for the study of the relations of the two other attributes manipulated in any given experiment.

To facilitate the analysis, we adopt a consistent convention in grouping and labeling the six measures. In each case we distinguish a primary dimension (labeled A), a secondary attribute (B) and the attribute of position (P). In the shape/color example introduced above, the A-attribute is shape and the B-attribute is color, by hypothesis.

The following examples assume that the training patterns are:

Pink circle on left

Blue triangle on right

- P? A+ This is the position question, when the A dimension votes 'yes' (the B-dimension votes 'no'). For example, 'IS POSITION NORMAL?' for a blue circle on the left.
- P? B+ Same question, with B-dimension voting 'yes', e.g., pink triangle on left.

Note that the comparison of the answers to the two questions provides a fairly direct measure of the relative weights of shape and color, or of the strength of the tendency to interpret the question 'is position normal?' as 'normal for this shape' or 'normal for this color'. The question can be construed as a passive version of the sorting task often used in categorization research: instead of asking subjects to assign events to a location, we provide a location and require them to evaluate it.

An unanticipated result of our study was that the P? question provides a sensitive diagnosis for a failure (or rejection) of categorization: subjects sometimes find the test pattern so distinctive that they do not spontaneously categorize it with either of the training patterns. The question 'Is position normal?' then appears odd, because no normal position has been established for that test pattern, and the common answer is 'No'. What makes this result interesting is that it does not always occur: although every test pattern is abnormal, many test patterns are said to appear in a normal position.

- P- A? Here we take one of the training patterns, present it on the 'wrong' side (e.g., blue triangle on the left) and query the primary attribute ('IS SHAPE NORMAL?')
- P- B? Same stimulus, with secondary attribute queried.

The two questions refer to a stimulus in which a norm display is simply shifted from one position to the other. The display attributes are normally coupled with each other, but both are in the wrong place. In answering either question, position presumably votes 'No', but the color is normal for the shape that is presented, and the shape is normal for the color. This analysis suggests that the two attributes of the object should not be judged equally normal: the dominant attribute could be abnormal because it is most directly tied to position, whereas the secondary attribute could be normal because it is paired in the usual fashion with the dominant one. The hypothesized pattern of responses to these two questions is our principal measure of contingent coding. The second major surprise of the study was that this expected pattern was quite elusive. In most cases subjects assigned similar ratings to the dominant and to the secondary attribute, even when their responses to other questions gave clear evidence of the differential importance of these attributes. Identifying the restricted conditions under which contingent coding does occur is the major focus of our current work with this technique.

- P+ A? Here the primary attribute appears in the same

position in the neighboring norm, but the secondary attribute belongs to the competing norm. For example, 'IS SHAPE NORMAL?' for a blue circle on the left.

P+ B? An example is 'IS COLOR NORMAL?' for a pink triangle on the left: it is normal for objects on the left to be pink, and normal for triangles to be blue. Which will prevail?

The answers to the P+ questions provide another measure of the relative weight of the two test attributes in determining each other's judged normality: in both cases the untested attribute votes 'no', whereas position votes 'yes'. Note that the stimuli for the P+A? and for the P+B? questions are respectively the same as the stimuli for the P?A+ and P?B+ questions. However, subjects quite often give a positive answer to one of the corresponding questions and a negative answer to the other.

The six normality questions provide three independent tests of the relative importance of the two critical dimensions: the dimensions are pitted against one another in the first set of questions, and their influence on one another is compared to that of position in the second and third sets. The original hypothesis was that the three tests would generally agree: the more important (less mutable) dimension was expected to control sorting or categorization, as indexed by the P? question. The facts turned out to be more complex than my simple notions of mutability and dimensional importance had suggested.

Results and Discussion

In the following section we present the main results of the experiments, grouped by topic.

Color vs. Shape

There is interest going back at least to Heidbreder (1948), and as recently as Landau, Smith and Jones (1988) in the role of different attributes in classification, and in the particular hypothesis that adults and even young children find it more natural to classify objects by shape than by other attributes. In part as an attempt to validate the technique, we investigated this question in a number of experiments. Table 2-1 summarizes the experiments in which the critical attributes were shape and color.

Table 2-1

P?		P-		P+		Condition
A+	B+	A?	B?	A?	B?	
85	24	50	52	79	53	simple shapes -- distinct colors
91	12	34	50	91	53	complex shapes -- distinct colors
87	02	54	79	87	48	simple shapes -- similar colors
67	07	57	57	90	83	rectangles varying in orientation
50	32	30	50	86	71	rectangles varying in length
58	29	33	29	87	92	rectangles varying in aspect ratio

The first row of Table 2-1 illustrates the dominance of shape in two variations in which the shapes were distinctive geometric figures and the colors were highly discriminable. Two of the three manifestations of dominance are present: subjects are very likely to judge position normal if the shape of the test pattern corresponds to the shape of the training pattern usually shown in the same place. In the third test (the P+ questions) subjects are also significantly more likely to judge abnormal a color that is paired with the 'wrong' shape than a shape paired with the wrong color. However, there is no indication of dominance in the P- questions: when a training pattern is shown intact in the wrong place, subjects tend to assign the same ratings of normality to both attributes: the mean judgments are similar and the correlation is substantial.

The next rows of the Table summarize several experiments designed to clarify the role of discriminability of both attributes. Making the colors quite similar (though still easily distinguishable) had little effect. On the other hand, the dominance of shape was clearly reduced when the shapes were rectangles -- they appeared in that case as color patches. The most important result in this series is the effect of varying shape by adding a small feature, as illustrated in Figure 2. This feature is less obvious and almost certainly less discriminable than the color difference between the norms, by standard measures, such as identification threshold or speed of same-different judgment. Although formal experiments will be needed to nail down the point, we are confident that neither discriminability nor an impression of within-attribute differences can explain the pattern of normality judgments. In general, the variations of within-attribute similarity, although they had some effect, did not reverse the dominance of shape over color. Along this line, Landau, Smith and Jones (1988) report that subjects over two years of age persist in classifying by shape even when a test stimulus is twelve times the size of a training stimulus.

Garner (1974) dealt with the same problems in his studies of free classification of objects varying on integral and separable dimensions. (Our measures of spontaneous categorization differ from his in that we do not explicitly require subjects to classify objects or events). Garner presented evidence that free classification is quite insensitive to manipulations of within-attribute differences, and he concluded that his classification data reflected genuine dimensional preferences (see also Garner, 1983). Shape and color are of course separable by Garner's criteria. His conclusion about the reality of dimensional preferences appears to hold for spontaneous categorization as well as for free classification.

Another question that should be raised concerns the relationship between normality judgments and similarity judgments. Could our

results be explained by assuming that subjects answered the normality question by grouping the test stimulus with the norm to which it was globally most similar, and answered the normality question accordingly. We have direct evidence that this account could not be completely correct: the data provide a number of instances in which the two normality judgments obtained for the same test display differed very markedly. This could hardly occur if the normality judgments merely reflected global similarity.

It is quite possible, however, that the dominance hierarchy of dimensions which we observed in normality judgments would also be observed with similarity judgments. We examined this possibility in a pilot experiment, in which subjects first made normality judgments, then similarity judgments in which they compared each stimulus to the two norms. One of the test cases was the pair of shapes shown in Figure 4. The two figures differ by a small extra feature of shape. They were made strikingly different in coloring by giving one shape a bright multicolored pattern, leaving the other a solid dull red. Nevertheless, most subjects (10 of 12) appeared to judge global similarity by shape rather than by color. We are not sure that this result would have been obtained without the earlier test with normality judgments, a point that will be clarified in further work. If the result is stable, however, the conclusion need not be that the similarity results 'explain' the normality judgments, because the similarity results are themselves quite puzzling. Garner (1974) argued long ago, that similarity judgments follow quite different rules when the stimuli differ on separable or on integral dimensions, and he was quite suspicious of similarity in the former case. Tversky and Gati (1982) also drew a sharp distinction between the two cases, pointing out that a geometric representation is possible for integral but not for separable dimensions. The issue of the relative weight, salience or prominence of various dimensions can be raised -- and should be raised -- for similarity as well as for normality (Shepard (1964) and Nosofsky (1986) have considered factors that affect the weights of dimensions within a metric model). We plan to explore this topic in detail in the next phase of the research.

The observations on shape and color collected so far suggest two conclusions: (1) The dominance of shape is a robust result, which may depend more on the individuality of the shapes than on their discriminability. (2) Color is subordinate to shape in these judgments, but is not nested within the shape attribute -- i.e., the judgment of the normality of color is not screened off from the attribute of position.

Shape vs Other Object Attributes

Experiments involving different combinations of static properties are summarized in Table 2-2. A dominant attribute is easily identified in each one of these experiments, except for the color/size comparison, which was ambiguous. The dominance of shape over other attributes was

confirmed in the answers to the P? and P+ questions. However, even more consistently than in Table 2-1, the judgments made in the P- condition uniformly fail to show the expected difference in the normality of the dominant and subordinate attributes.

Table 2-2*

P?		P-		P+		Condition
A+	B+	A?	B?	A?	B?	
90	10	67	73	80	33	shape/size
93	10	60	67	90	33	shape+color/size
70	13	53	63	97	50	shape/number
46	52	52	72	76	59	color/size
65	22	57	52	59	35	Number/color

The clear dominance of shape over number is a noteworthy result; it surprised us because the global appearance of the display appeared to us to be determined by the number of stimuli (three or four, arranged respectively in a fairly widely spaced triangle or parallelogram), in that experiment).

Object Pairs

Many of our experimental conditions involved a pairing of distinct objects, instead of attributes. For example, the two training norms might consist of (1) a large pink L with a green bar over it, presented on the left of the screen; or (2) a large pink T with a green star over it, presented on the left. The test stimuli would recombine the pink and green objects in the two positions in the six possible ways. Our general hypothesis in this phase of the study was that the salience of the two objects would determine their weight in the spontaneous categorization process which is (we hoped) tapped by our questions. Table 2-3 summarizes experiments in which the patterns were defined by combinations of parts or by sequences of events rather than by conjunctions of attributes.

Table 2-3

P?		P-		P+		Condition	
A+	B+	A?	B?	A?	B?		

Position							
34	31	28	25	56	78	high shape (A), low shape	
31	34	59	47	78	66	left shape (A), right shape	
34	37	16	31	75	81	low letter (A), high letter	
44	25	50	31	81	53	left letter (A), right letter	
09	19	53	50	87	94	front shape (A)/back shape	
Codability and position: letters and shapes							
41	53	44	31	87	44	letter below shape	
66	47	59	56	62	37	letter above shape	
25	25	50	44	78	59	letter to left of shape	
47	28	31	37	97	53	letter to right of shape	
Size							
40	27	50	73	97	70	large (A) below small (B)	
Size and codability							
*	75	00	11	39	86	71	large letter, small shape
	54	25	12	17	87	83	large shape, small letter
Integration of large and small parts							
	37	37	33	54	87	79	connected parts, ask by 'shape'
	50	25	25	43	93	79	connected parts, ask by 'part'
	59	19	62	75	37	31	arm of machine, ask by arm, object
*	83	13	60	60	90	47	small and large parts of Gestalt
	42	08	37	29	87	62	hats next to faces
Embedding and circling							
	36	31	67	71	57	45	configuration of small shapes
	31	19	53	53	78	75	stationary frame around object
	31	21	50	67	93	78	object moving from a frame
	25	19	51	68	82	81	object moving around another

The main result in Table 2-3 is negative: it appears to be much more difficult to obtain dominance in combinations of parts than in conjunctions of attributes. In some conditions involving framing of a shape, for example, each training pattern consisted of a central shape or character surrounded by eight characters of the same size. When the frame and the center were recombined in test patterns, subjects usually judged the position abnormal, and evaluated the normality of the constituents by their relationship to position rather than to each other.

The tendency of subjects to judge position abnormal for any recombination of parts is observed in many of the experiments. The sum of values for P?A+ and P?B+ is .75 or less for 17 of the 23 values shown in Table 2-3. Negative answers to the position question are common even when one of the parts is clearly more salient than the other (i.e., when there is a substantial difference between the evaluation of the normality of position for the A+ and B+ cases). What makes this result interesting is the contrast with the results observed for conjunctions of shape and other attributes, and for objects and actions. For comparison, the answers to the two position questions add up to less than .75 in only 2 of 33 experiments reported in Tables 1, 2, 4 and 5 in this section. We have interpreted the position question as a test of spontaneous sorting, or categorization. By this test, subjects appear less willing to categorize by constituent elements than by properties. The generally high ratings of normality for the less salient constituent in the P+ condition indicates that it was usually coded in terms of its position on the screen, not in relation to its more salient neighbor.

The two most successful attempts to obtain a dominance effect in the P? questions are marked by stars in Table 2-3. Clear dominance was achieved by pairing a large letter with a small shape (the size ratio was 3:1 in height and width). Dominance was also observed with the two norms shown in Figure 3, which invite the coding of a main object with a part that fits into it (or fails to fit).

Sequences

Table 2-4 summarizes experiments in which the display consisted of a series of events. In the first of these experiments a letter was seen, which appeared to move to another location and simultaneously to change into another letter. We had expected that the second event in the sequence would be coded as subordinated to the first, and perhaps even as contingent on it. Nothing of the kind happened. The pattern of results for that condition was very similar to that observed with compound displays (see Table 2-3). The failure to obtain dominance in this simple situation is a significant result, because it eliminates a plausible interpretation of normality ratings as reflecting the confirmation or disconfirmation of expectations. The critical comparison is between the ratings of the first and of the second events in the P+ condition. In the A? case the first event

occurs in its usual place and is followed by an unusual sequel; in the B? case, the second event does not correspond to the letter that just preceded it. However, subjects rated both events normal, indicating that the relation to position was more important than the sequence of expectations and confirmations.

Table 2-4

	P?		P-		P+		Condition
	A+	B+	A?	B?	A?	B?	
Sequences							
	27	27	58	42	96	92	simple sequence of two letters
*	69	31	81	77	69	50	character appearing in frame
*	60	17	60	63	53	37	sequence of two distinctive motions
*	72	03	31	34	66	56	complex motion --> simple motion

Some dominance was evident in the three other types of sequences. These were slightly more complex than the simple sequence of letters, in an attempt to induce conditionalization on the first part of the sequence. In the first, a character appeared in the center of a frame of letters. The last two were sequences of different motions, one matched for complexity of motion, the other presented a complex motion followed by a simple one.

Objects and Motions: Conditions for Contingent Coding

Dominance of one attribute over another was quite often observed in the results presented so far, but contingent coding was striking by its absence. Contingent coding was defined in the section that introduced the normality technique by a particular pattern of answers on the P- questions: low on the A? question, higher on the B? question. Evidence of contingent coding is finally found where there was most reason to expect it, in judgments of the normality of objects and their actions. The precise conditions under which contingent coding is found -- in contrast to mere dominance -- now appear to be quite an interesting problem, which we plan to explore in further work.

Table 2-5 presents results for conditions involving objects, actions and changes. We briefly consider the conditions in turn.

Something resembling the expected pattern of contingent coding was obtained where the norms were schematic faces of a boy and a girl, one frowning and the other smiling. Phrasing the normality questions was awkward. We ended up asking "Is the character normal?" for the A? question and "Is the expression normal?" for the B?. Half of the responses to the former question in the P- condition were positive,

indicating a tendency to judge the norm display as normal when it is presented intact in a new position. The proportion of positive responses to the B? question was very significantly higher, suggestive of contingent coding. However, the character was judged abnormal when paired with the wrong expression in the P+A? condition.

Table 2-5

P?		P-		P+		Condition
A+	B+	A?	B?	A?	B?	
97	10	47	87	27	17	faces and expressions
89	26	46	80	61	30	motion, ask by object
83	39	57	72	48	28	occluded, ask by object
96	23	62	92	38	35	motion (frozen), ask by object
77	31	54	88	46	31	occluded (frozen), ask by object
83	83	25	33	92	83	motion/color, ask by color
92	75	29	29	96	71	occluded/final color, ask by color
68	54	43	50	82	50	shape/motion, ask by shape
71	75	39	46	82	54	occluded, shape, ask by shape
72	72	31	44	97	66	shape/motion, large shapes that sit
87	84	37	56	81	59	shape/occluded, large shapes
77	53	40	70	87	30	shape+color/motion, ask by shape
80	53	53	73	100	37	shape+color/occluded, ask by shape
98	07	39	87	52	20	destination, ask by object
96	31	42	81	31	19	destinations (frozen), by object
93	43	29	71	93	32	shape/destination, ask by shape
93	27	47	77	93	20	shape+color/destinations, by shape
96	14	39	71	93	68	changing colors
87	03	37	56	91	50	shape/changing salient colors

A rather similar set of results was obtained with small objects in simple motions (e.g., a blue star moving up as the left-side norm and a yellow cross moving down as the right side norm). The A? and B? questions referred respectively to the normality of the object and of the motion. In the 'occluded' variant the object moved either up-then-down, or down-then-up; the moving object was a white diamond on the first leg of the motion, which turned into a colored shape at that point, yielding an impression that the 'true' moving object was then revealed. Note that in this condition the observer identify the motion before they can identify the moving object.

The judgments for the regular and the occluded variants of motion are generally quite similar. This is a significant observation, which shows again that the normality ratings are not determined by the sequence of confirmed or disconfirmed expectations. Regardless of the order in which the information is received, the P? results confirm the expected dominance of the object on the motion. The answers to the A?P+ question indicate that the subjects tended to describe an object as abnormal if it behaved in the manner associated with the other object. Finally, and most importantly for the present discussion, a substantial minority of subjects said that the object was abnormal but the motion normal in the P- condition. The results were essentially the same when the objects simply disappeared at the end of their motion or when they remained frozen in their terminal positions.

Additional experiments were carried out in an attempt to identify the conditions that produced the new pattern of results in the motion and occluded-motion experiments. As shown in the Table, one feature of the results depends on asking subjects about the normality of the object, rather than of a particular attribute: low ratings were given in the P+A? condition, presumably because the object appeared abnormal when it behaved abnormally. This was not true when respondents evaluated the normality of 'the shape'. The more significant feature of the results is the discrepancy between the normality ratings for shape and motion in the P- condition: the dissociation vanished when the moving objects were distinguished by a single feature (either shape or color); it was restored when the objects were defined by a conjunction of features, although the normality of the shape (not the object) was judged. I return later to a discussion of this observation, on which our current work is focused.

A robust pattern of contingent coding was also obtained in another condition, in which objects appearing on one or the other side of the display (labeled 'starting positions' in the questions) moved to one or another marked destinations, above and below the center of the display (labeled destinations). In the P- condition the object (or 'shape', or 'color' in different experiments) is consistently judged less normal than the destination. Evidently, the evaluation of the destination is conditioned on the object, not on its starting position.

Some evidence of contingent coding is also present in a 'changing color' display, in which the norms are distinctive white shapes, which gradually take on distinctive colors. The situation is informationally equivalent to that investigated in the shape/color experiments described earlier, but the judgments indicate a stronger tendency to rate the shape by its position and to relate the color to the shape.

Verbal Label as a Medium

In the experiments mentioned so far position was the medium used to study the interactions between other attributes. As noted earlier, we interpreted the responses to the P? questions as equivalent to sorting the test stimuli with one or the other of the two norms. This interpretation gains credence from the frequent cases in which the answers to the two P? questions added up close to 100%, suggesting that the subjects evaluated the normality of position entirely by its correspondence with one of the other attributes.

The use of a verbal label as a medium is an obvious alternative to position. Distinct labels can be attached to the two norms (now shown in the same location), and the two values of this attribute (called LABEL in the questions presented to the subject) then replace the values of position in the standard design. We initially chose position over label as a medium because we wished to evoke immediate perceptual judgment rather than a reflective problem-solving attitude, which may be likely to be applied to a more verbal task. Having collected most of our data with position, however, we checked a selected set of conditions in an experiment using label as a medium.

The procedure was identical to the experiments using position as a medium, except that instead of stimuli appearing on different sides of the monitor, they appeared in the center. A three letter nonsense syllable, either "aba" or "uzu", appeared on the left side of the monitor with every trial.

The results are shown in Table 2-6.

Table 2-6

L?		L-		L+		Condition
A+	B+	A?	B?	A?	B?	
66	16	12	19	87	81	shape/color , small feature
62	09	16	44	91	81	shape/color, hard to name shapes
56	37	34	37	91	69	destinations
56	34	31	31	91	87	connected parts, ask by 'part'
62	06	31	34	91	75	large letter, small shape
50	34	41	19	94	66	motion
53	28	28	34	87	87	occluded motion
56	25	22	16	84	78	sequence of complex-simple motion

The attributes that dominated the answer to the P? questions in earlier experiments are consistently assigned the greater weight in judgments of the normality of labels, but the sum of answers to the L? questions is generally rather low. Subjects would evidently much prefer to retain the label for the conjunction of properties with which it was associated. This preference appears to be stronger for labels than for positions.

Another salient feature of Table 2-6 that the attributes of the various displays are judged very largely in terms of label: the normality of the secondary attribute in the L+B? condition is generally very high; the incongruous value of the dominant attribute has little effect. Thus, the label appears to be so dominant that it swamps other effects. This feature of the results vindicates the original choice of position as a medium.

Discussion

Three main patterns of results have appeared in these studies of normality judgments.

(1) In compound displays subjects tended to relate each object directly to position rather than to the other object in the display. They also tended to judge position abnormal if either of the constituents appeared in the wrong place. Extreme manipulations of codability and size were required to produce dominance of one of the constituents in the encoding of the scene. These negative results indicate that the normality judgments do not simply reflect differences in salience or physical similarity.

(2) A pattern of dominance of shape over other attributes was observed in many experiments. Correspondence to shape controlled the judgments of the normality of position (P?), and there often was an effect of shape on the judged normality of the attribute with which it was paired (B?P+). However, contingent coding was not found: subjects

appeared to provide a single judgment of normality to the intact display presented in an unusual position. Although formal control experiments are yet to be run, it is quite clear that the differential weight of shape and other attributes in normality judgments does not depend on corresponding differences in the discriminability or physical salience of these attributes: shape dominated size and number, and small but distinctive differences in shape dominated color. Distinctiveness rather than discriminability appears to be the key to dimensional dominance.

(3) At still higher levels of object distinctiveness, and particularly when change or movement are involved, the pattern of contingent coding is observed, with object dominance in the P? and P+ conditions and substantial dissociations of normality judgments in the P- condition. Our results suggest that a conjunction of small differences may be more effective in producing this pattern than a highly discriminable difference in one attribute. The notion of distinctiveness that comes up here recalls some classic writings by George Miller (1956) and Wendell Garner (1965, 1974). In his famous discussion of the magic number seven, Miller pointed out that multiple attributes are necessary to permit large number of entities to be discriminated from each other in memory -- quite a different problem from sensory discrimination between two simultaneous or immediately successive stimuli. Garner (1965, 1974) offered the idea that stimuli define the set of their alternatives. In the present experimental situation the differences between the two norms define the set of alternatives. This set is much larger when the norms differ a little in several attributes than when they differ greatly in one attribute. The distinctiveness or individuality of the stimuli increases with the size of the pool of possibilities from which they have been drawn. The present observations on the weighting of conjunctions are not isolated: Medin and Edelson (1988) observed that conjunctions were given greater weight than single features in a categorization task, with diagnosticity controlled.

Part 3. Topic and Referent in Symmetric Comparisons (Research conducted by Maria Stone)

Comparison is an important aspect of the encoding and interpretation of experience. Comparison is also a skill that is executable under voluntary control: when instructed, one can choose a subject or topic, and compare it with a standard or referent, on specified dimensions or in a general search for differences. There is usually more than one way to express the result of a comparison, e.g., "A is larger than B" or "B is smaller than A". For many relationships, however, one of the directions is used more often and appears more natural to native speakers (D'Arcais, 1970; Talmy, 1984; Jackendoff, 1988).

There is a compelling intuition that the distinction between topic and referent does not only arise at the stage of choosing a verbal

description for a comparison; the topic and the referent are processed differently in the comparison itself. In the context of our general search for hierarchical and asymmetric relations in coding, we became interested in the factors that determine the spontaneous allocation of the roles of topic and referent when these roles are not prescribed by instructions. An attractive hypothesis arises from a list of the characteristics that control the choice of the subject of a comparative or locative sentence (Talmy, 1987). Talmy's list includes such characteristics as moving as opposed to stationary, recent on the visual scene, smaller, simpler, more salient, etc. The list overlaps substantially with the list of factors that control the spontaneous allocation of attention to objects in a scene (Berlyne, 1950, 1951, 1958; Kahneman, 1973). Maria Stone has carried out several experiments to examine the role of attention in the allocation of roles in comparisons. We were also concerned with comparisons between a test item and a previously presented standard: are there conditions under which the nominal standard is actually the topic of comparison?

A thought experiment (which we have carried out as an actual experiment on a few subjects) will suggest the strength of the effect of attention on verbal encoding. Imagine that you are shown a red line and a green line on a screen. One of the lines is stationary, the other is moving. A sentence to be verified appears on the screen, e.g., "it is longer than the red line". It is intuitively obvious (and obvious in our preliminary data) that the sentence will be checked more quickly if the moving line is green than if it is red.

Our interest, however, is to isolate the cognitive effect of attention on topic selection from its effect on verbal encoding. To this end, Maria Stone adapted a method used by Agostinelli et al (1988), which exploits an asymmetry noted by Tversky (1977): the similarity of A to B is reduced by features that are unique to the topic (A) more than by features that are unique to the referent (B). Similarity data or same/different judgments can therefore be used to diagnose which of the two elements of a comparison served as topic; the difference between two stimuli will be noted more readily if the unique features that distinguish them belong to the topic rather than to the referent. To determine how the order of presentation affects the direction of comparison, Agostinelli et al. (1986) constructed pairs of pictures, consisting of a simple picture and a more complex one, derived by adding pictorial elements. Subjects encountered one picture as a study item, and could encounter the other as a test item, in a subsequent recognition task. This design allows a comparison of accuracy with which difference is recognized when the item presented at test includes new features, or lacks features that were present in the corresponding study item. Agostinelli et al concluded that the test item is likely to be selected as the topic of comparison, unless the subjects have advance knowledge of the subsequent task and a chance to prepare for it.

Our experiments employ a similar logic to determine the role of attention in comparison. We require subjects to make same/different

judgments about pairs of stimuli. When the stimuli are different, the difference is produced by adding unique features to one of them.

Method:

The stimuli were 20 pairs of modified Chinese characters. Each pair consisted of a basic character and a complex character, generating by adding two extra lines. Each character was drawn in a grid 4 cm. wide and 4 cm. long. Viewing distance was 50 cm. The characters were drawn and displayed in red on a black background, using our standard color graphics system.

The task was a same/different discrimination between two characters. After one block of practice, subjects had five blocks of 80 trials.

On each trial two characters were simultaneously presented. There were eight possible positions, equally spaced on a circle (radius 4 cm.). The two characters shown on a trial appeared at the two ends of a diameter. One of them was stationary, the other in apparent horizontal oscillation of 4mm. extent. The stationary phase and the ISI were each 80 msec.

Results:

TABLE 3-1: RESULTS OF EXPERIMENT 1
MOVING/STATIONARY

response attended object	same		different	
	complex	simple	complex	simple
reaction time	1269	1077	1003	1109
error rates(%)	2.0	2.1	1.1	4.7

The results are shown in Table 3-1. The hypothesis of the study was strongly confirmed: the same difference between two stimuli was more rapidly detected when the unique features that distinguished them belonged to the moving (presumably attended) shape than to the other. The effect is observed on reaction times ($t(9)=6.74$, $p<.0005$) and on error rates ($t(9)=2.88$, $p<.01$). We conclude that the moving character was selected to be the topic of comparison, with the stationary as a referent. Note that the effect is obtained without contamination by linguistic factors: there was nothing in the situation to encourage subjects to encode the comparison verbally.

The experiment also provides a control for a possible compatibility effect that had complicated the interpretation of our first pilot studies on the problem. In these studies we presented a red line and

a green line, and asked subjects to press a key corresponding to the longer line.

Attention was manipulated by oscillating one of the lines. As expected, the response was faster when the attended line was longer, and also when it was presented in the fovea. There is no artifact of discriminability in these studies, because the subject must use information from both lines to make a decision. However, our initial experiments did not eliminate the possibility that attending to a line of a particular color evoked a tendency to make the response corresponding to that line. The use of the same/different paradigm precludes this possibility.

Experiments 2, 3, and 4.

Many comparisons are made between an object presented at test and an object presented earlier -- and presumably retrieved from memory for the test. What is the topic in such comparisons? Agostinelli et al found cases in which the test item appeared to be the topic, and other situations in which the item retrieved from memory appeared to be compared to the test item. As noted earlier, the test item was the topic in an unexpected recognition task at the end of the experimental session. In another experiment, the study and test items were presented in immediate succession. A detailed analysis suggested that the study item served as topic for codable stimuli (i.e., line drawings of real recognizable objects--cars, houses, etc.), but not for uncodable ones (collections of geometric shapes and lines). We carried several experiments to follow up this work, using pseudo-Chinese characters varying in complexity. The results available at the time of writing are described below.

Experiment 2 was designed to create a situation in which we would be essentially certain that the topic of comparison -- if any is selected -- would be the stimulus presented most recently. To this end, we first presented two characters, then added to the display a third character that was to be compared to both, for a same/different judgment. In Experiments 3 and 4 the two characters to be compared were presented sequentially; the experiments differed in the presence of a mask following the first character, and in ISI.

Method

The stimuli for the three experiments were chosen from the same set of 20 pairs of Chinese characters. In the first part of each trial two characters were shown above the fixation point. After 960 msec a third character was added below the fixation point. The new character either matched one of the old ones exactly or was the simpler/more complex version of one of the old characters. The subjects indicated the presence or absence of an exact match by pressing a key.

In Experiment 3 one character was shown for 480 msec, followed by a mask of 480 msec and a blank ISI of 960 msec, prior to the

presentation of the next element of the comparison. In Experiment 4 a single character was presented for 480 msec, followed with ISI = 480 msec by a second character.

Results and Discussion

The results of these experiments are summarized in Tables 3-2 - 3-4. As expected, a character added to an existing display of two characters became the topic of the comparison, using the diagnostic measure employed in Experiment 1: the 'different' RT was substantially faster when the new stimulus was complex than when it was simple, ($t(7)=4.1$, $p<.005$), and the difference in error rates is also significant and in the same direction ($t(7)=3.00$, $p<.01$).

TABLE 3-2: RESULTS OF EXPERIMENT 2
RECENCY ON THE VISUAL SCENE

response	same		different	
	complex	simple	complex	simple
first object				
reaction time	1124	908	1003	935
error rates(%)	5.1	3.3	4.9	7.1

An equally clear result was obtained in Experiment 3. With a mask and a long ISI intervening between the two characters, the first stimulus presented was selected as topic. This result was observed on reaction times ($t(8)=4.08$, $p<.005$), as well as on error rates ($t(8)=3.29$, $p<.01$).

TABLE 3-3: RESULTS OF EXPERIMENT 3
FIRST-SECOND WITH MASK AND ISI=960 MSEC.

response	same		different	
	complex	simple	complex	simple
first object				
reaction time	800	791	810	882
error rates(%)	11.3	13.0	14.5	28.1

Finally, Experiment 4 did not indicate any significant tendency to select one or the other stimulus as topic on reaction times ($t(7)=.03$), though accuracy was significantly better when the complex character appeared first ($t(7)=2.67$, $p<.05$). The difference between the results of Experiments 3 and 4 is highly significant ($t(15)=2.98$, $p<.01$).

TABLE 3-4: RESULTS OF EXPERIMENT 4
FIRST-SECOND WITH ISI=480 MSEC

response first object	same		different	
	complex	simple	complex	simple
reaction time	701	636	739	739
error rates(%)	8.6	9.3	10.0	14.2

These results indicate that there is no general rule for the stimulus that is present at the time of test to be the topic of a comparison. A stimulus that is no longer present can be allocated that role, and compared to the item currently available. However, we suspect that this strategy requires a form of active rehearsal of the physically absent item. If confirmed, this hypothesis would unite the analyses of the effects of order of presentation and of attention on topic selection. A rehearsed item presumably is the focus of attention, and therefore a favored candidate for the role of topic. This hypothesis is readily testable, by using visual filler tasks to suppress rehearsal.

Part 4: The Language of Counterfactuals

In preparation for a conference on counterfactuals and simulation, organized by Jim Sherman in April 1989, Carol Varey and I began an informal analysis of the language of counterfactuals, with a view to learn from language about the psychology of counterfactual thinking. Specifically, we considered the felicity conditions for the phrases 'X came close to happening' and 'X almost happened'. The exercise proved quite productive, and Sherman suggested that this unusual effort be written up and submitted to the Journal of Personality and Social Psychology, which he edits. The writing is in progress. In this section I briefly summarize the major points.

One thing we can learn from language is that counterfactual assertions refer to the objective world, not a state of mind. Thus one can say "He came close to dying but no one knew it". The counterfactual part of the sentence specifies something about the world, not about anyone's belief at the time. Another indication is that counterfactuals can be faked "I used to think that professional wrestlers often come close to killing one another, but in fact they do not". Thus, the statement that "Seton Hall almost defeated Michigan in the final" is a statement about the game, not about the observers' state of mind.

'Coming close to happening' involves a causal process that unfolds over time. It expresses more than mere propensity. Thus, a die may be loaded to show six 80% of the time. When the die shows three,

however, we cannot say that it came close to showing six. The latter assertion requires an activated propensity, a causal process that has been initiated.

The evaluation of a counterfactual works backward, from the point in which the 'factual' was realized. The tracing back sequence is an essential aspect of causal thinking, which was emphasized by Hart and Honore (1959) in their classic Causation and the Law. As Hart and Honore pointed out, tracing back tends to stop at intentions or else at 'causal joints'. Consider the examples

- 1) They came close to marrying, but she never intended to do it.
- 2) She came close to accepting a Harvard offer, but it didn't come.

Both 1 and 2 are anomalous, because one-sided willingness to enter into an agreement is not sufficient for 'coming close'. Now consider what the speaker would need to know in order to justify the following statements:

- 3) Frank came close to robbing a large bank.

- 4) Bob came close to stealing money from his mother.

Note that Frank's desire to money and an absence of moral restraints do not suffice for the assertion that he came close to robbing a bank. To be appropriate, such an assertion must rest on knowledge about Frank's criminal skills, about detailed plans he has made, possible partners he has recruited, and so on. The situation is quite different when it comes to stealing from most mothers: in the absence of significant obstacles to the success of the enterprise, serious consideration of the crime may suffice for coming close. The entire causal structure in which the factual occurrence is embedded is pertinent to judgments of counterfactual alternatives.

To appreciate the role of causal joints, imagine a party that went out on two pleasure boats. People chose a boat at random. One boat sank with no survivors. A person who was on the other boat can rightly say "I came close to going on the doomed boat", but cannot say without elaboration "I came close to drowning".

Now consider a baby carriage that rolls down an incline, passes some distance from a tree and stops unharmed. What is the maximal distance at which it is still reasonable to say that 'the carriage almost hit the tree'? An important variable is whether the carriage had a baby in it. The statement would be acceptable at a larger distance for an occupied carriage, suggesting a role for emotional focus. Indeed, the acceptable distance is larger if the speaker is the baby's mother than if it is a stranger. Intricate pragmatic calculations are evidently involved.

The role of emotional focus is evident in many other examples, e.g., "I almost died" vs "he almost died", "we almost scored another touchdown" vs "they almost scored another touchdown" etc. Two distinct interpretations are possible. A cognitive interpretation is that focusing of attention may increase mutability (Kahneman & Miller, 1986). A pragmatic interpretation is that some things that are worth

saying about an occupied carriage are not worth saying about an empty one. Separating the two interpretations raises an interesting experimental challenge.

Finally, we analyze rather subtle judgments about a more abstract concept of distance. Consider the statements "they almost bought the house" and "they almost did not buy the house". Which of the two statements allow more latitude -- a greater distance from realization? The question is initially bewildering, but a large majority of our casual collection of linguistic informants agree that it is appropriate to say that "they almost bought the house" if they considered it seriously; it takes much more to justify saying "they almost did not buy". In the latter case, listeners are led to infer that the deal was very close to unraveling at some point.




The asymmetry of allowed distance in 'almost' is quite a general phenomenon. The main determinant appears to be the identity of the outcome that is considered normal in the relevant context. In the absence of context information, a default is assumed. For example, buying a particular house is less normal than not buying it. In general, there is greater latitude in saying that an abnormal occurrence almost happened than there is in the reverse case.

A rather simple-minded linguistic analysis appears able to shed light on quite interesting cognitive phenomena, including the representation of counterfactual alternatives, the processing of causal information, and the pragmatic coordination of norms. We are currently extending this analysis, and are also planning some more systematic studies to assess the uncontaminated intuitions of non-specialists.

References



- Agostinelli, G., Sherman, S. J., Fazio, R. H., & Hearst, E. S. (1986). Detecting and identifying change: Additions vs. deletions. Journal of Experimental Psychology: Human Perception and Performance, 12, 445-454.
- Ashby, G. and Perrin, N. (1988). Toward a unified theory of similarity and recognition. Psychological Review, 95(1), 124-150.
- Bartlett, J. C. and Dowling, W. J. (1988). Scale structure and similarity of melodies. Music Perception, 5(3), 285-314.
- Berlyne, D. E. (1960). Conflict, arousal, and curiosity. New York: McGraw Hill.
- Berlyne, D. E. (1950). Stimulus intensity and attention in relation to learning theory. Quarterly Journal of Experimental Psychology, 2, 71-75.
- Berlyne, D. E. (1951). Attention to change. British Journal of Psychology, 42, 269-278.
- Berlyne, D. E. (1958). The influence of complexity and novelty in visual figures on orienting responses. Journal of Experimental Psychology, 55, 289-296.
- Chase, W. G. and Clark, H. H. (1971). Semantics in the perception of verticality. British Journal of Psychology, 62(3), 311-326.
- Chase, W. G. and Clark, H. H. (1972). Mental operations in the comparison of sentences and pictures. In Gregg, L.W. (Ed.) Cognition in Learning and Memory, New York: N.Y.: John Wiley & Sons, VII.
- Clark, H. H. & Clark, E. V. (1978). Universals, relativity, and language processing. In J. H. Greenberg, C. A. Ferguson, & E. A. Moravcsik (Eds.), Universals of Human Language, (pp. 225-277). Stanford, CA: Stanford University Press.
- Clark, H. H., Carpenter, P. A. and Just, M. A. (1973). On the meeting of semantics and perception. In Chase (Ed.) Visual Information Processing, New York, N.Y.: Academic Press, 311-379.
- Clark, H. H. and Chase, W. G. (1972). On the process of comparing sentences against pictures. Cognitive Psychology, 3(3), 472-517.
- D'Arcais, F. G. B. (1970). Linguistic Structure and focus of comparison in processing comparative sentences In D'Arcais and Levelt (Eds.) Advances in Psycholinguistics, Amsterdam-London: North-Holland Publishing Company, 307-321.
- Erbaugh, M. S. (1986). Taking stock: the development of Chinese noun classifiers historically and in young children. In C. Craig (Ed.), Typological studies in language: Vol. 7. Noun classes and categorization (pp. 399-436). Amsterdam: John Benjamins.
- Garner, W. R. (1962). Uncertainty and structure as psychological concepts. New York: Wiley.
- Garner, W. R. (1974). The processing of information and structure. New York: Wiley.
- Garner, W. R. (1983). Asymmetric interactions of stimulus dimensions in perceptual information processing. In T. J. Tighe & B. E. Shepp (Eds.), Perception, cognition and development: Interactional analyses. Hillsdale, NJ: Lawrence Erlbaum.

- Gentner, D. & French, I. M. (in press). In S. L. Small, G. W. Cottrell, & M. K. Tanenhaus (Eds.), Lexical Ambiguity Resolution, Los Altos, CA: Morgan Kaufman.
- Hart, H. L. A. & Honore, A. M. (1959). Causation in the Law. London:Oxford University Press.
- Heidbreder, E. (1948). The attainment of concepts: VI. Exploratory experiments on conceptualization at perceptual levels. Journal of Psychology, 26, 193-216.
- Heidbreder, E. (1949). The attainment of concepts: VII. Conceptual achievements during card-sorting. Journal of Psychology, 27, 3-39.
- Heidbreder, E., Bensley, M. L., & Ivy, M. (1948). The attainment of concepts: IV. Regularities and levels. Journal of Psychology, 25, 299-329.
- Heidbreder, E. & Overstreet, P. (1948). The attainment of concepts: V. Critical features and contexts. Journal of Psychology, 26, 45-69.
- Jackendoff, R. (1987). The relation of linguistic and visual information. Cognition, 26, 89-114.
- Just, M. A. and Carpenter, P. A. (1976). Eye fixations and cognitive processes. Cognitive Psychology, 8, 441-480.
- Kahneman, D. & Miller, D. T. (1986). Norm theory: comparing reality to its alternatives. Psychological Review, 93, 136-153.
- Kahneman, D. (1973). Attention and Effort, Englewood Cliffs, N.J.: Prentice-Hall, Inc.
- Landau, B., Smith, L. B., & Jones, S. S. (1988). The importance of shape in early lexical learning. Cognitive Development, 3, 299-321.
- Medin, D. L. & Edelson, S. M. (1988). Problem structure and the use of base-rate information from experience. Journal of Experimental Psychology: general, 117, 68-85.
- Miller, G. A. (1956). The magical number seven plus or minus two: some limits on our capacity for processing information. Psychological Review, 63, 81-97.
- Mithun, M. (1986). The convergence of noun classification systems. In C. Craig (Ed.), Typological studies in language: Vol. 7. Noun classes and categorization (pp. 379-397). Amsterdam: John Benjamins.
- Nosofsky, R. M. (1986). Attention, similarity, and the identification-categorization relationship. Journal of Experimental Psychology: General, 115, 39-57.
- Patterson, E. (1930). A qualitative and quantitative study of the emotion of surprise. Psychological Monographs, 40, 85-108.
- Shepard, R. N. (1964). Attention and the metric structure of the stimulus space. Journal of Mathematical Psychology, 1, 54-87.
- Talmy, L. (1987). The relation of grammar to cognition. In Rudzka-Ostin and Brygida (Eds.) Topics in Cognitive Linguistics, John Benjamins, forthcoming.
- Talmy, L. (1983). How language structures space. In Pick and Acredolo (Eds.) Spatial Orientation: theory, research and application, New York: Plenum, 225-282.
- Tversky, A. (1977). Features of similarity. Psychological Review, 84, 327-352.
- Tversky, A. and Gati, I. (1982). Similarity, separability, and the triangle inequality. Psychological Review, 89, 123-154.





IS POSITION NORMAL?

IS COLOR NORMAL?



IS POSITION NORMAL?

IS SHAPE NORMAL?



IS COLOR NORMAL?

IS SHAPE NORMAL?




Figure 1. Example of an experiment.





Figure 2. Shapes differing by a small feature.





Figure 3. Stimuli that form a good Gestalt.





Figure 4. Small difference in shape, large difference in color.